



Law Society
of Ontario

Barreau
de l'Ontario

TAB 5

AI for Litigators

The Path of the Law: Toward Legal Singularity

Professor Benjamin Alarie, Osler Chair in Business Law, Faculty of Law
University of Toronto

April 16, 2021



The Path of the Law: Toward Legal Singularity

May 27, 2016

Benjamin Alarie, Osler Chair in Business Law
Faculty of Law, University of Toronto

ABSTRACT: In 1897 in the *Harvard Law Review*, Oliver Wendell Holmes Jr. wrote that, “For the rational study of the law the blackletter man may be the man of the present, but the man of the future is the man of statistics and the master of economics.”¹ Now, nearly 120 years later, the future that Holmes foresaw is arriving. Much of the global population has transitioned from an analog, paper-based world with unreliable, slow and costly communication, to a digitally connected world with almost universal real-time and nearly costless communication. And within this digital world we are witnessing substantially greater availability of data and improved methods of machine learning through advances in computer-assisted modelling and inference. The implications for law of an abundance of data of all kinds and dramatically more effective statistical tools are becoming visible. The ultimate consequences for law will be profound. I propose referring to the culmination of these developments as “the legal singularity.” The legal singularity will affect all areas of the law. For the purposes of illustration, I focus my attention here on tax law. I predict that coming decades will witness three gradual transitions as the legal singularity draws nearer: (1) improved dispute resolution and access to justice in tax law, primarily through the transition from our current reliance on standards (adjudicated *ex post*) to greater reliance on query-able systems of complex rules (knowable *ex ante*); (2) a transition to superior and increasingly more complete specifications of tax law (*i.e.*, a gradual transition from the complex, unwieldy, uncoordinated tax systems of today to tax systems that are massively complex and yet precisely and effectively distribute benefits and burdens); and, (3) with the realization of the legal singularity, a complete specification of tax law (and, indeed, all the other areas of law), which will thenceforth remain (more or less) in positive and normative equilibrium. The equilibrium achieved by the legal singularity will be a type of reflective equilibrium along the lines described by John Rawls in *A Theory of Justice*.²

Acknowledgements: I am grateful to Anita Anand, Richard Bird, Tony Casey, Daniel Debow, David Dyzenhaus, Mohammad Fadel, Andrew Green, Birke Hacker, Brett Janssen, Edward Iacobucci, Karen Knop, Kai Konrad, Anthony Niblett, Arthur Ripstein, Erik Roder, Geordie Rose, Emily Satterthwaite, Wolfgang Schoen, Hamish Stewart, Katie Sykes, Albert Yoon, and to workshop participants at the University of Toronto and at the Max Planck Institute for Tax Law and Public Finance, for comments, questions, and stimulating conversations surrounding these ideas. All responsibility for errors rests with me alone.

¹ Oliver Wendell Holmes Jr., “The Path of the Law” (1897) 10(8) *Harvard Law Review* 457, at 469.

² See John Rawls, *A Theory of Justice* (Harvard University Press, 1971) at 48-53.

1. Introduction

Lawyers and judges have mental and moral abilities that (so far) surpass those of machines.³ Human comparative advantages lie in the parsimony of our reasoning and our ability to reach serviceably good conclusions with woefully incomplete information. Our law schools teach us to read cases, to identify the most important principles from those cases, and then to construct conjectures and theories about how novel fact patterns might be addressed by judges in future disputes. We accomplish this by reasoning by analogy to different areas of the law, by contemplating the relationship between moral principles and legal principles, by drawing upon our intuitions, and so on. Many commentators, particularly before the advent of modern computing, registered deep reservations about the possibility of machines ever being able to convincingly emulate these human abilities (let alone surpass us in these tasks),⁴ while more recently others took a rosier view.⁵

And yet the presence of commentators expressing reservations about the plausibility of various advances of technology, including in the legal context, is hardly new.⁶ Richard Susskind's prediction in *The Future of Law* in 1996 that email would become a tool commonly used by lawyers drew criticism and skepticism.⁷ Commentators committed to the status quo have been frequently discredited by our experience of progress. Past decades are littered with examples of machines exceeding our greatest human champions at tasks that for time immemorial seemed beyond the reach of mere machines. The world chess champion, Garry Kasparov, fell to Deep Blue in 1997. The most winningest human players of 'Jeopardy!' ever, Ken Jennings and Brad Rutter, fell to IBM's Watson in 2011. Google's AlphaGo defeated Lee Sedol, one of the

³ A majority of experts predict that by year 2100 that our machines will have developed general artificial intelligence. See Vincent C. Müller and Nick Bostrom, "Future Progress in Artificial Intelligence: A Survey of Expert Opinion" in Vincent C. Müller (ed.), *Fundamental Issues of Artificial Intelligence* (Synthese Library; Berlin: Springer), 553-571.

⁴ See, for example, B.C. Brosnahan, "The Law and Computers" (1970) 1(3) *Auckland U. L. Rev.* 1 at 2: "no matter how many contingencies the program may be designed to cope with, a computer can never outdo humans and perform activities that cannot be analysed into logical patterns."

⁵ See, for example, Daniel Martin Katz, "The MIT School of Law? A Perspective on Legal Education in the 21st Century" (2014) 2014 *Illinois Law Review* 1431; Mitchell E. Kowalski, *Avoiding Extinction: Reimagining Legal Services for the 21st Century* (American Bar Association, 2012); and Richard Susskind and Daniel Susskind, *The Future of the Professions: How Technology Will Transform the Work of Human Experts* (Oxford University Press, 2015).

⁶ See, for example, JC Smith, "Machine Intelligence and Legal Reasoning" (1998) 73 *Chicago-Kent Law Review* 277, at 311: "The reality of legal reasoning is far too complex to describe and comprehensively represent" and at 345: "There is only one kind of intelligence, and that is human."

⁷ Richard Susskind, *The Future of Law: Facing the Challenges of Information Technology* (Oxford University Press, 1996), at 242.

game's strongest human players, at the game of Go in March 2016.⁸ Chess and Go had each, in past decades, been declared to be the ultimate litmus test of the effectiveness of machine learning and artificial intelligence.

Perhaps the naysayers are correct in claiming that humans will never be exceeded by machines in the context of law. Excelling at law ostensibly requires a sound heart as well as an able mind. It is unclear whether our machines will be able to emulate human empathy and moral sentiments and, even if not, whether these are in fact required for legal acumen. Despite general uncertainty about the specifics of the path ahead for the law and legal institutions and what might be required of our machines to make important contributions to the law, over the course of this century we can be confident that technological development will lead to (1) a significantly greater quantification of observable phenomena in the world ("more data"); and (2) more accurate pattern recognition using new technologies and methods ("better inference"). In this contribution, I argue that the naysayers will continue to be correct until they are, inevitably, demonstrated empirically to be incorrect. The culmination of these trends will be what I shall term the "legal singularity."⁹

The legal singularity will arrive when the accumulation of massively more data and dramatically improved methods of inference make legal uncertainty obsolete. The legal singularity contemplates complete law. The legal singularity is inspired by and different from the idea of the technological singularity popularized by the futurist Ray Kurzweil. The technological singularity refers to the stage when machines themselves become capable of building ever more capable and powerful machines, to the point of an intelligence explosion that exceeds human understanding or capacity to control (technological singularity is akin, then, to superintelligence).¹⁰ The legal singularity contemplates the elimination of legal uncertainty and the emergence of a seamless legal order, universally accessible in real-time. In the legal singularity, disputes over the legal significance of agreed facts will be rare. They may be disputes over facts, but the once found, the facts will map on to clear legal consequences. The law will be functionally complete.¹¹

The legal singularity will affect all areas of the law. For the purposes of illustration, I focus my attention here on tax law. I predict that coming decades will witness three gradual transitions

⁸ See Christof Koch, "How the Computer Beat the Go Master" (March 19, 2016) *Scientific American*, <http://www.scientificamerican.com/article/how-the-computer-beat-the-go-master/> <accessed April 18, 2016>.

⁹ The idea of a coming technological singularity was popularized by Ray Kurzweil, *The Singularity is Near: When Humans Transcend Biology* (Penguin, 2006).

¹⁰ On the implications of technological singularity and superintelligence, see Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press, 2014).

¹¹ The two ideas are connected in the sense that should a Kurzweil technological singularity come to pass, it would facilitate the legal singularity through the sheer sophistication of data collection and machine inference thereby made possible. Because the two ideas have distinct meanings, a technological singularity is not a necessary condition for the legal singularity.

as the legal singularity draws nearer: (1) improved dispute resolution and access to justice in tax law, primarily through the transition from our current reliance on standards (adjudicated *ex post*) to greater reliance on query-able systems of complex rules (knowable *ex ante*); (2) a transition to superior and increasingly more complete specifications of tax law (*i.e.*, a gradual transition from the complex, unwieldy, uncoordinated tax systems of today to tax systems that are massively complex and yet precisely and effectively distribute benefits and burdens); and, with the realization of the legal singularity, (3) a complete specification of tax law (and, indeed, all the other areas of law), which will thenceforth remain (more or less) in positive and normative equilibrium. I address each of these transitions in parts 2 through 4, below. Part 5 concludes.

2. From Standards to Complex Query-able Systems of Rules

Standards in the law provide useful rules of thumb to guide human behaviour. The first year law student who struggles to apply legal standards in individual legal cases (*e.g.*, whether a defendant in a negligence case “acted reasonably”) nevertheless has generated from his or her human experience a general capacity to imagine or understand what it means to “behave reasonably” in a particular human social context. This is true even if the law student would be uncomfortable attempting to prospectively identify the complete set of actions that would constitute reasonable behaviour in a given setting. The usefulness of such generalities as “behaving reasonably” is clear: these standards accord with our intuitions, meet our expectations for the conduct of others, and allow us to proceed more or less confidently in the world. Essentially, standards permit judges to decide after the fact what legal result or results ought to obtain, all things considered and with the benefit of hindsight. Legal standards provide useful heuristics about what is legally permissible. In the future of tax law, we will increasingly witness legal standards becoming superseded by complex systems of rules that perform the same function as legal standards but with the important advantages of greater reliability, predictability, and accessibility in real time.¹²

Let’s consider how standards are starting to be cast into complex query-able systems of rules in Canadian income tax law. In tax law, standards govern many issues that are fundamental to the determination of tax liability. One frequently arising issue is whether an individual worker is better classified as an employee or as an independent contractor. The distinction between the two categories is clear in the paradigmatic cases. On the one hand, for example, the plumber who agrees to a fixed fee on a one-time-only basis to replace a bathroom fixture in an individual’s private home is, *vis-à-vis* the hiring homeowner, clearly an independent contractor. On the other hand, a plumbing supply company with plumbers on payroll, who must wear company-provided uniforms, have company-defined holiday schedules, enjoy pension entitlements, and who perform tasks on behalf of the company from 9am to 5pm, Monday to Friday, use the company’s vehicles and tools, and are barred from working for any other

¹² For an excellent extended treatment of a related idea, see Anthony Casey and Anthony Niblett, “The Death of Rules and Standards” (November 20, 2015) *U of Chicago, Public Law Working Paper No. 550*; *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 738*; available at SSRN: <http://ssrn.com/abstract=2693826>.

plumbing supply companies, are employees. The challenge in tax law arises in characterizing relationships that fall in the relatively uncharted territory between these extremes.

Canadian tax law relies on the common law to classify workers as independent contractors or employees.¹³ After hundreds upon hundreds of cases, judges have been unable to articulate a single bright line test that separates independent contractors from employees.¹⁴ Instead the courts use something that the judges refer to as the “total relationship” test.¹⁵ The “total relationship” test looks to all aspects of the relationship to determine what is the most appropriate classification of the worker and identifies certain aspects that should specifically be investigated: the intention of the parties, who exercises control over the worker, which party has the ownership of tools, the worker’s risk of loss and chance of profit, and the integration of the worker into the hirer’s business activities. The courts do not prescribe the weights that ought to apply to these various criteria, and are clear that these criteria are not necessarily exhaustive; other considerations might apply in an appropriate case.

A fascinating thing is that (in one sense at least) there *is* a precise test that is available—it is implicit in the hundreds of judgments decided over the past two decades. Analyzing these hundreds of cases using new machine learning techniques suggests that there is a reasonably clear (but extraordinarily complex) distinction that can be drawn from the decided cases. Even at this relatively early stage of applying machine learning to this body of case law, research with two of my colleagues shows that it is possible to provide dramatically improved classifications of workers *ex ante*—that is, at the time that workers are being hired, without having to wait to learn what the tax authorities or the judges of the Tax Court of Canada might have to say about the matter.¹⁶ Using a hand-collected dataset of all of the cases from the Tax Court of Canada involving worker classification from the 1990s through to early 2016, and using computationally intensive machine learning algorithms, it is now possible to predict the classification that would be given to a worker with greater than 98% confidence on the basis of answers to just twenty simple questions relating to the relationship between a worker and a hirer.¹⁷

Why does this matter? In the Canadian income tax law system, and in other tax systems too, taxpayers often have an incentive to try to achieve one characterization rather than another. In

¹³ In Quebec, the situation is somewhat different. For an overview of the legal issues and the case law, see David G. Duff, “The Federal Income Tax Act and Private Law in Canada: Complementarity, Dissociation, and Canadian Bijuralism” (2003), 51(1) *Canadian Tax Journal* 1, at 26-32.

¹⁴ For background, see Alain Gaucher, “A Worker’s Status as Employee or Independent Contractor” in *Report of Proceedings of the Fifty-First Tax Conference, 1999 Conference Report* (Toronto: Canadian Tax Foundation, 2000), 33:1-98.

¹⁵ See, for example, Joanne E. Magee, “Whose Business Is It? Employees Versus Independent Contractors” (1997) 45(3) *Canadian Tax Journal* 584-603.

¹⁶ Research conducted by Benjamin Alarie, Anthony Niblett and Albert Yoon from the University of Toronto, Faculty of Law.

¹⁷ *Ibid.*

the context of the classification of workers, hirers frequently attempt to establish relationships with workers so that the workers are classified as independent contractors rather than as employees. The reasons for this preference include streamlined administration—including not having to withhold and remit income taxes and payroll tax at source (required of employees but not independent contractors); not having to pay the employer’s portion of payroll taxes; not having to run payroll, but rather to pay invoices; and allowing the workers to claim for themselves a broader array of deductions as independent contractors than they would be permitted as employees. These tax considerations often lead parties to put into place arrangements that they expect will reasonably lead to a characterization of a worker as an independent contractor rather than as an employee. Sometimes the arrangements are unambiguously successful, in which case the tax authorities are not bothered, or at least they should not be bothered. Sometimes, however, the parties proceed less carefully and claim the benefits of an independent contractor relationship when it is not justified.

The tax consequences of the total relationship test can be significant. Tax authorities are often in the position of pursuing those taxpayers whom it feels have not classified their workers properly, and taxpayers are often in the position of having to defend the claim that their workers are independent contractors. These disputes arise frequently at the Tax Court of Canada. The court releases dozens of judgments each year. The volume of litigation suggests that taxpayers are not clear about how the total relationship test will be applied in individual cases. Indeed, because details surrounding the entire relationship are relevant to the characterization, in ambiguous circumstances it is difficult to predict how a court would resolve this question. What is clear is that many taxpayers, undoubtedly aware of the risk of misclassification, nevertheless conclude that they are better off taking the position that their workers are, indeed, independent contractors, rather than employees. By using a classification tool like those currently in development, hirers can set the terms of their relationships with their workers in a way that allows them to be confident about what the classification would be if the case were to be decided by the Tax Court of Canada.

The benefits extend beyond resolving the taxpayers’ legal uncertainty. Tax authorities can be more confident about their ability to consistently and fairly administer differences of opinion involving the classification of workers on audit. Judges, too, can avail themselves of these new tools to ensure that they are efficiently taking into account all of the previously decided case law that applies to similar cases. In the context of worker classification, more data and better inference tools make possible much sharper predictions about the content of standards in the law. If one is unclear about how to classify a worker, a taxpayer merely needs to consult with the applicable tool and an extremely reliable answer can be provided in minutes.

It should be clear that this sort of approach extends to many other questions in tax law, for example, whether a particular expenditure is more accurately characterized as being current or capital, or whether the sale of an asset should be considered to be on income account or on account of capital. These sorts of classifications are sometimes easy to make, but in the grey

areas become difficult.¹⁸ They are perhaps less difficult judgments to make, however, than many judges and lawyers imagine once one subjects the entirety of the case law to careful scrutiny and trains a machine learning system to identify the patterns that lead to the outcomes in the cases.

One of the reasons why machine learning systems may find these kinds of judgments easier to make than human lawyers and judges is that the machine learning systems face different constraints and limitations than individual human beings. There is a large literature in psychology that documents the various limitations that humans confront in reasoning and memory tasks.¹⁹ Machines are not as intellectually versatile as humans (at least not yet). One advantage that machines do have over humans is the size and reliability of their short- and long-term memories, as well as their ability to carry out strings of logical operations. A machine trained on the distinction between employees and independent contractors can simultaneously take into account every judgment that it has been exposed to in training, can construct elaborate models of relationships among the variables accounted for in the case law, and can quantify with precision the result of the algorithms used to evaluate the possibilities.

With the considerable advantages that machines have over humans in terms of memory, objectivity, and logic, one may feel that machines will in the near future come to strictly dominate humans in the law. Although there is considerable evidence that experts expect this to be true in the long run, for the next several decades (in all likelihood) it is likely that more data and better machine learning inference tools are likely to be complements to human judgment rather than substitutes. A survey of experts in the field project that there is a 90% chance of artificial general intelligence being developed by 2075.²⁰ In the meantime, a clever human lawyer or judge will know when and under what circumstances the conclusions offered up by a classification prediction system are likely to be suspect, and when they are likely to be more reliable. For this reason, it is likely that for several decades a human expert when paired with big data and a machine learning system will be superior to either kind of expert considered alone.²¹

What developments like these mean for tax law is that compliance with detailed and complex tax law will be easier than ever before. Various standards of tax law (and, indeed, other areas of the law), once reduced to query-able systems of complex rules, will enable an initial vetting of

¹⁸ For example, according to the Master of the Rolls, Sir Wilfred Greene in *British Salmson Aero Engines Ltd v. CIR* (1938), 22 TC 29, at 43: "... there have been ... many cases where this matter of capital or income has been debated. ... in many cases it is almost true to say that the spin of a coin would decide the matter almost as satisfactorily as an attempt to find reasons." Quoted in *Johns-Manville Canada v. The Queen*, [1985] 2 S.C.R. 46 at para. 13.

¹⁹ For an introduction on many of these topics, see Robert A. Wilson and Frank C. Keil, *The MIT Encyclopedia of the Cognitive Sciences* (MIT Press, 2001).

²⁰ See Müller and Bostrom, *supra* note 3.

²¹ An example of this can be seen in the game of chess: the strongest current "players" are teams made up of strong human players using chess software on fast hardware.

plans at low cost. Professional advisors will continue to play a role in helping taxpayers navigate the tax system, particularly for those with complicated situations or a desire to engage in more adventurous means of avoiding tax liability. Advisors working with machine learning systems will be able to devise ever more complex and ingenious tax avoidance plans that strain the capacity of the tax administration to respond effectively. For this reason, governments (including the courts, the tax authorities, and policymakers) will use the machines as well. The result will be a legal arms race, with the machines playing a central role in the escalation of tensions.

3. Toward Legal Singularity: Closing the Gaps in the Law

There is already widespread dissatisfaction with tax law and policy throughout developed countries. There is considerable tax avoidance, particularly among the most sophisticated taxpayers, multinational corporations and wealthy individuals and families.²² In the shadow of current technology and political arrangements, the administration of tax law depends upon voluntary self-assessment. The precepts of the rule of law require that the law be announced and promulgated in advance, which gives taxpayers and their clever advisors lead time to identify and exploit the incompleteness of the law. The incompleteness of the law leads to over- and under-inclusiveness. There are circumstances in which enforcing the tax law as written may have undesirable consequences and lead to unanticipated double-taxation. Moreover, even if satisfied that the tax law was perfectly specified and would remain so, tax administrations in all countries simply do not have the information, administrative capacity, and enforcement power to strictly enforce tax law to the letter.

Currently, the tax law is incompletely specified, in a manner that is suboptimal. Supporting evidence for this view is the presence of general anti-avoidance and anti-abuse provisions in many of the world's leading income tax and value-added tax systems, some of which have been introduced relatively recently (e.g., the general anti-abuse rule in the UK, which was introduced in 2013).²³ Optimal setting and enforcement of the tax law would do away with the need for general anti-avoidance rules and, for that reason, would require massive amounts of data. The sort of legal understanding that would be demanded to completely specify tax law would include domestic and international law, as well as the interaction of all underlying concepts in tax, private law, commercial law, family law, etc., across all jurisdictions, be they subnational, national, and supra-national. Given limited policymaking resources and the reality of political economy, it is not surprising that tax law is not now complete, nor perfectly specified.

Nevertheless, governments will come to increasingly rely on machine learning to identify weaknesses in the system and to implement safeguards against tax avoidance. They will have little choice as a consequence of the ever more elaborate tax plans advanced by advisors with

²² See Benjamin Alarie, "The Challenge of Tax Avoidance for Social Justice in Taxation" in H.P. Gaisbauer et al. (eds.), *Philosophical Explorations of Justice and Taxation* (Springer International, 2015), 83-98.

²³ UK HMRC, "Tax Avoidance: general anti-abuse rule" (last accessed April 18, 2016) <https://www.gov.uk/government/publications/tax-avoidance-general-anti-abuse-rules>.

access to query-able systems of complex rules. Drawing upon these machine learning inspired anti-avoidance measures, the tax system will become even more complex. The complexity will increase with time such that eventually everyone will become dependent on machine learning to cope with the complexity of the system. Interestingly, a by-product of this dynamic is that despite the ever-increasing complexity of the system, the effect of the law will be more predictable and reliable than ever. Tax law will become even more complete and, as gaps and loopholes in the law are closed, will become more completely specified.

One of the likely consequences of increasingly relying on big data and user-friendly systems in order to make decisions in tax is an acceleration of the dynamic back and forth between governments seeking to collect revenues and taxpayers avoiding taxes. To understand the intuition underlying how we can expect this to play out, consider the above example of worker classification. If tax accountants and tax lawyers consult a machine learning tax classification tool in advising taxpayers about how to precisely arrange their affairs so that it is 99%+ likely that their workers will be characterized legally as independent contractors, the tax authorities and the government might well be concerned that the tax base will be eroded; taxpayers will, for their part, worry that they are being forced to jump through unnecessary hoops to secure a tax result that they may have helped themselves to with little fear of discovery in years past. Before long, tax policy makers will likely suggest certain changes to the law in order to defend (or restore) the tax base and tax revenues, or to facilitate more efficient workings of the system (e.g., consider the “check-the-box” regulations in the US for entity classification). One way might be to reduce the tax advantages of classification as independent contractor rather than employee (that is, by increasing the neutrality of the tax system). Another way would be to be more specific and dictate more elaborate criteria for the employee-independent contractor distinction (by increasing the difficulty of manipulating features of a worker-hirer relationship to secure the desired classification). Either way, the law will be become more complete by being better specified.²⁴

With time, machine learning systems will continue to improve and obviate the necessity of having a human tax advisor as intermediary, especially for taxpayers with relatively common circumstances. The machines themselves will be able to fully “understand” the law and the taxpayer’s factual context sufficiently well to make authoritative suggestions about how to optimize the taxpayer’s behaviours and plans. In the long run, so long as a taxpayer conveys his or her values and preferences to the system, the taxpayer will (correctly) be able to trust that the machine learning system has optimized its advice based on the law and the facts. Governments, for their part, will use machine learning systems to further optimize the content of the tax law, given prevailing politically-endorsed social values. Accurate modeling of the mutually optimal strategic responses will be based on vast amounts of data on taxpayer behaviour, generated by the abundant data made available by widespread sensors (in currently

²⁴ This idea is similar to theories of evolution of the common law, though the selection mechanism of evolution in this context is different. For an example of an empirical analysis of the evolution of the common law, see Anthony Niblett, Richard A. Posner, Andrei Shleifer, “The Evolution of a Legal Rule” (2010) 39(2) *Journal of Legal Studies* 325.

fashionable parlance, the “Internet of Things”). Politics and policymaking, domestic and international, will be dominated by rich debates surrounding what values tax systems ought to embrace and reflect. Machine learning systems assist in implementing tax and transfer systems that achieve the distributive justice trade-offs that democratic political processes endorse.

4. The Legal Singularity: A New Positive and Normative Equilibrium

Where more data and better inference take us may well be to the realm of what would currently be regarded as science fiction. It could be that we collectively begin to rely on big data and machine learning to assist us in formulating normativity. In part 2, the discussion suggested that big data and machine learning can help us to complete the law by substituting extremely complex, fact-sensitive, and query-able systems of rules for existing legal standards. In other words, machines can help us to identify what the law should be in order to achieve our implicit social objectives. In this part, the argument is that big data and machine learning can go a step further and assist us in formulating the objectives that we collectively should embrace and embody in our legal systems.

As computing power increases, it is conceivable that machine learning systems will develop the capacity to explore the normative dimensions of taxation (and other public policies), with an eye to bridging the gap between normative values and politics and lived factual consequences. For example, it seems possible (if not likely) that machine learning systems will be able to identify a menu of optimized policy packages to select from. This would represent, in effect, an “efficient frontier” of policy tradeoffs. And if it is plausible to imagine a machine learning system identifying an efficient frontier in terms of public policy, it is perhaps not a leap too far to speculate that a machine learning system might be able to identify a uniquely “best” tax policy package, given the specified normative objective. In the long run, it appears likely that machine learning systems will debate amongst themselves and alongside our most accomplished and esteemed philosophers, economists, artists, and scientists.

Should the evolution of more data and better inference ever reach this point, we should expect the resulting reflective equilibrium to be relatively stable.²⁵ Until now, public policy has been made with such limited data and imperfect information that it should be no surprise that the law is constantly changing. In the legal singularity, the equilibrium that would be reached would be based on positive and normative convergence of facts and values. Although I may be willing to concede that the parameters of the resulting legal system might well oscillate back and forth in as circumstances change and responses to various economic shocks (e.g., the emergence of new technologies) might merit short-term fiscal accommodation, we should not anticipate that there would be significant developments leading to marked changes in the legal system. If achieved, the legal singularity will generate a stable and predictable legal system whose oscillations will be continuous and yet relatively insignificant.

²⁵ See the discussion in Rawls, *supra* note 2 at 49-50. It should be clear that I am more inclined to believe that the sort of reflective equilibrium that Rawls described might be plausibly achieved, though not by unaided human philosophers alone.

5. Conclusion

More data and machine learning will generate powerful tools for the improvement of the legal system. Ultimately, I believe these developments will result in the “legal singularity” which results in a more or less positively and normatively stable legal system. The apotheosis of the legal system will be extraordinarily complex and will be beyond the complete understanding of any person.

One of the major developments as tax law becomes better specified and as technology permits will be that the tax system will increasingly be used to deliver benefits as well as distribute burdens. At the moment, of course, tax systems are already thought to principally be in the business of distributing burdens and ensuring that all taxpayers contribute a fair share to help support government. In future, it is likely that much of the economic production will be carried on by non-humans. In this environment, governments may come to rely on the taxation of those with significant means derived from non-human production of economic value to help those humans who do not work. One possibility is that the tax system will be used to deliver a universal basic income, which would be provided to all taxpayers, and then treated as income and taxed back at higher levels of income.

These predictions turn on two major assumptions. The first is that technological progress continues to generate more data. The second is that our methods for analyzing data continue to improve due to increases in computing power and better methods of machine learning. It may be objected that there is a third assumption: that taxpayers and governments alike will welcome (or at least not openly resist) the emerging role of machine learning into the tax system through the law, the facts, and values. This assumption is, of course, implicit and by no means guaranteed. What grounds are there for expecting cooperation or, at least, no staunch resistance that pre-empts these developments? At least for the next several decades, it appears likely that tax professionals will be indispensable and valuable complements to machine learning. In the long run, sometime beyond 2050, we will probably face the need to do a radical rethinking of a number of tax-related roles, including: judges, lawyers, accountants, academics, etc. These roles will be probably preserved in some altered form, but reshaped dramatically by the effects of technological development.